arXiv:1607.07931v1 [cs.CL] 27 Jul 2016

Synthetic Language Generation and Model Validation in BEAST2

Stuart Bradley
Supervised by Dr. David Welch

July 28, 2016

Abstract

Generating synthetic languages aids in the testing and validation of future computational linguistic models and methods. This thesis extends the BEAST2 phylogenetic framework to add linguistic sequence generation under multiple models. The new plugin is then used to test the effects of the phenomena of word borrowing on the inference process under two widely used phylolinguistic models.

Contents

Acknowledgments **1**

1 Introduction **2**

2 Models and Algorithms for Simulating Language Evolution **4**
2.1 Cognates . . . 4
2.2 Evolutionary Trees . . . 4
2.3 Cognate Substitution Models . . . 5
2.3.1 General Time-Reversible (GTR) Model . . . 6
2.3.2 Covarion Model . . . 8
2.3.3 Stochastic-Dollo Model . . . 9

3 Whole Tree Sequence Generation **12**
3.1 Algorithm Validation . . . 13

4 Borrowing **14**
4.1 Global Borrowing Model . . . 14
4.2 Local Borrowing Model . . . 15
4.3 Adding Borrowing to Cognate Substitution Models . . . 15
4.3.1 GTR Borrowing . . . 15
4.3.2 Stochastic-Dollo Borrowing . . . 19

5 No Empty Meaning Class **22**

6 Missing Data **22**
6.1 Missing Languages . . . 23
6.2 Missing Meaning Classes . . . 24
6.3 Algorithm Validation . . . 25

7 Computational Optimizations **25**

8 Model Misspecification **27**
8.1 Testing for Misspecification . . . 27
8.2 Misspecification Due To Borrowing . . . 27
8.2.1 Material and Methods . . . 28
8.2.2 GTR Model . . . 31
8.3 Stochastic-Dollo Model . . . 32
8.3.1 GTR Inference . . . 33
8.3.2 Covarion Inference . . . 35
8.4 Discussion . . . 36
8.4.1 Topology . . . 37
8.4.2 Tree Height . . . 37

9 Conclusion **38**

A Language Sequence Generation Plugin Guide **42**
A.1 Quick Start 42
A.1.1 Command Line 42
A.1.2 BeastBorrowingPluginTest 42
A.2 Explanation of the Input/Output Files 43
A.2.1 The BEAST File 43
A.2.2 The Output File 45
A.3 The Thesis Analysis Package 45

Acknowledgments

If I begin by saying this thesis wasn't a serious trial, I would be doing both you (as the reader), and myself a major disservice. Despite my many trepidations towards doing postgraduate academic work, there are a sizable number of people that dragged me through honours — at times kicking and screaming. It is these next few paragraphs that I dedicate to them, and all their help and hard work.

Firstly and most importantly, I must thank my supervisor: Dr. David Welch. While his expertise and guidance were greatly appreciated, I am most thankful for his patience in dealing with not only a stubborn, but at times an ignorant student. Thank you for sticking with me, and helping me create something I am truly proud of.

Next up I would like to thank members of the newly created Centre for Computational Evolution. Not only for allowing me to bounce around ideas and thoughts, but also for aiding me in both the theoretical and computational pursuits of this thesis. While the entire department was of great help, I'm going to use the following lines to list a number of instrumental people: Dr Remco Bouckaert, Dr Luke Maurits, Dr Walter Xie, William Hsu, Megan Ganley, and Arman Bilge. Thank you to you all.

Outside of the group I'd also like to give thanks to Prof Clark Thomborson and Dr Rizwan Asghar, who gave me advice on a previous paper, and helped me become far more confident in my academic communication. Additional thanks goes to Dr. Wayne Mitchell and Asela Dassanayake at Lanzatech for giving me the leeway required to focus on the thesis.

Finally, I come to my family (Jack, Alyson and Dana Bradley), my flatmates (Simon Corkindale, Kieran Newburgh, and Kathleen Seddon), and my large group of friends (with special reference to Erika Horlor). You have all been of immense help this past year, and I truly could not have done it without you.

Once again, thank you all so much, and I hope you (the reader) find some level of enjoyment and education in the following body of work.

1 Introduction

Historical linguistics attempts reconstruct the complete history of language families in a systematic and sound way. Over the years there have been number of approaches to doing this, yet there does not appear to be a single perfect method. Without such a method, any suggested approach must be rigorously tested and validated before attempting to stand up to serious scrutiny.

Comparative linguistics was the first such approach, helped in part by an explosion of linguistic data between 1785 and 1820 [Collinge et al., 1995]. Compiled word lists were compared as proxies for entire languages, and could be used to determine the similarity of two languages (through *cognate words*). While this method was by all intents systematic, it was often subjective, and lacked systematic attempts to validate the resulting hypotheses.

This subjectivity was partially quelled by the introduction of *lexicostatistics* [Swadesh, 1955]. Swadesh developed a list of 200 meanings that was considered universal among languages, known as a *basic vocabulary*. By gathering word lists associated with each meaning for different languages, languages could be compared on a level playing field, and in theory, produced more accurate relationships, than could be determined when using earlier subjective judgments.

Glottochronology is an extension of the basic vocabulary concept, that includes word replacement at a constant rate, under a *glottoclock*, related to the idea of a biological *molecular clock*. From this, the percentages of shared cognates between languages can be converted into the length of time separating the languages in question. Larger percentages then suggest more recent separation. However, there are a number of issues that make glottochronology controversial [Bergsland and Vogt, 1962], such as the requirement in traditional glottochronology of a stable clock rate. This has in turn lead to modified versions of glottochronology that attempt to fix earlier issues [Kruskal et al., 1973].

In biology, *phylogenetic inference* examines evolutionary relationships between organisms, which produces a phylogeny. This is done via the comparison of heritable traits such as DNA or morphological traits. Phylogenetic inference also encompasses the use of computational techniques to explore the space of probable trees–something that isn't touched upon in older comparative methods.

Phylogenetics also has the added benefit of pre-existing theory brought over from biology. When languages are represented by a string of absence or presence data for a number of traits, they can be treated in a similar manner to DNA sequences, so that very little modification is required to convert a biological phylogenetic method to a phylolinguistic one.

An example of this is in the selection of the best tree. As there are already a number of methods for finding biological evolutionary trees, the options for finding linguistic trees are also numerous [see Sleator, 2011, for a recent review]:

1. *Distance-Matrix* - Sequence information is used to calculate distances between all pairs of organisms. The resulting distance matrix is used to construct a tree. Trees with distances along branches between pairs of taxa which are closest to the distances in the matrix are considered the best.

2. *Maximum-Parsimony* - The best tree is the one that requires the smallest number of evolutionary events to occur.

3. *Maximum-Likelihood* - Assigns probabilities to all possible phylogenetic trees via a model of substitution. The best tree is considered to be the one with the highest likelihood score.

4. *Bayesian-Inference* - Prior probabilities, along with a likelihood function, are used to estimate the posterior density over possible phylogenies. Rather than a single best tree, Bayesian methods typically produce a sampled set of trees representing a range of possible scenarios. This is often done using Markov chain Monte Carlo sampling.

Initial attempts to enfold linguistics inside phylogenetic frameworks include work by [Gray and Jordan, 2000], [Gray, 2003], and [Atkinson et al., 2005]. In all cases, the new inference method was used test already established hypotheses. These new methods brought some rigour to long-running debates, which lead to testing new ideas within this framework.

Atkinson et al. [2005] used phylogenetic inference to look at the origin of the Indo-European family of languages. Bouckaert et al. [2012] addressed the same problem with the explicit incorporation of geography and produced strong evidence for Indo-European having its origin in the agricultural expansion out of Anatolia.

To make sometimes complex phylogenetic tools more accessible, a number of packages were developed. BEAST [Drummond et al., 2012] (and later BEAST2 [Bouckaert et al., 2014b]) were created. These programs were designed around using Markov chain Monte Carlo as the basis for Bayesian inference to construct phylogenies and test evolutionary hypotheses.

While BEAST2 has a sequence generating capacity, it is ill-suited for models of language evolution; currently lacking for example, common models used in the field such as the stochastic-Dollo (SD) model.

This thesis aims to extend the language data simulation capacity of BEAST2, and to use simulations to test the robustness of common phylolinguistic estimation techniques to model misspecification, including various types of borrowing and missing data. In doing so, a language synthesis package (LangSeqGen) will be described that can be used quickly and effectively in future experiments.

2 Models and Algorithms for Simulating Language Evolution

Before diving directly into the simulation, both the language data, and the tree-like way in which it evolves, need to be discussed; this then leads naturally onto the models used.

2.1 Cognates

When one language diverges from another, both meanings and the forms of individual words can change. However, change is often an incredibly long process that takes place over many separate small changes [Greenhill et al., 2010]. Cognates are words in different languages which share a common ancestry, and are expertly defined as having similar meanings. In biological terms, cognate words are synonymous with homologous traits. An example of one such cognate might be: *night* (English), *nuit* (French), and *noche* (Spanish), which all derive from the common root: *nók^{w}ts (Proto-Indo-European).

While defining cognates is a systematic process, there are certain non-tree-like forces–such as borrowing–that muddy the process. To combat this, the Swadesh word list was developed. The list is a collection of 100-200 meanings that chosen to be culturally universal [Swadesh, 1955].

Given a common core of universal words, it then possible to compare different languages on the basis the presence or absence of specific cognates. However, direct comparisons–in the form of Glottochronology–rarely hold [Bergsland and Vogt, 1962].

Instead of using cognate words for direct tree construction, they can be used as basis for defining a concise representation of a language, which can then be used in the models discussed below.

2.2 Evolutionary Trees

Evolution is often represented as a tree [Darwin, 1859]. Starting with an ancestral species represented by a single lineage, evolutionary divergence causes the tree to branch. Splitting may be due to various circumstances, such as the separation between populations, which creates two distinct groups. Each lineage then continues to evolve separately through time causing the species to diverge. At the tips (leaves) of the tree, are the most modern–and possibly extinct (due to lack surviving information)–species.

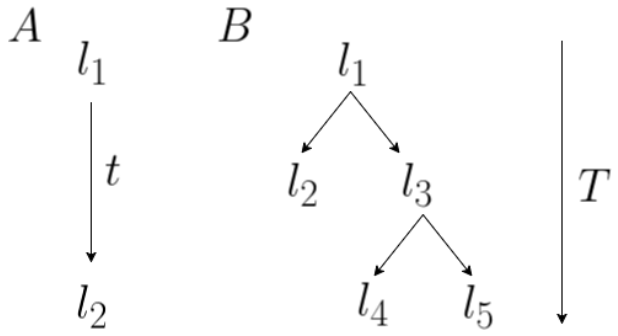

Figure 1: A: Single lineage evolution. B: Whole tree evolution.

Languages can also be represented in this fashion. In Figure 1 A, language l_2 diverges from its ancestral language l_1; as the result of multiple accumulated evolutionary events, over time t. Over a larger time frame T, this event happens many times, to multiple languages, which results in a language evolution tree (see Figure 1 B). In both cases, branch lengths can also vary, which is not shown in the figure.

This tree begins with a root language, which branches so that languages can diverge as they accumulate different mutations. Computationally, this divergence is actually represented as an instant event. Instead of slow divergence, l_1 is copied to l_2, and both languages then mutate in accordance with some model of language evolution. This effectively demonstrates the idea that l_2 slowly moves away from its parent l_1.

2.3 Cognate Substitution Models

When modeling language evolution, languages need to be represented as abstract forms so they can be modelled and processed efficiently. One approach to this is represent languages as the presence or absence of certain common cognate words, such as those in the Swadesh word list. For any language l, it can be represented as binary string. For example, if l were being represented by the presence or absence of four cognates, it might look like:

$$l = \begin{bmatrix} 0 & 1 & 1 & 0 \end{bmatrix} \tag{1}$$

Multiple languages can then be housed in a matrix (also known as an alignment), where each row represents a language and each column a cognate, to increase the compactness of the representation.

$$C = \begin{bmatrix} 0 & 1 & 1 & 0 \\ 0 & 0 & 1 & 0 \\ 1 & 1 & 1 & 0 \end{bmatrix} \tag{2}$$

Once a representation is specified, the next step is to specify a model of language evolution. The evolution of a language happens as the result of mutations which in this case are the gain and loss of cognates.

In the models discussed below, a site-independent *Markov process* is used as the basis for evolutionary effects on languages. Site-independent means that it is assumed that mutations at any site are independent from any other site. The Markov assumption implies that times between mutations are exponentially distributed according to a rate parameter defined by the specific substitution model.

2.3.1 General Time-Reversible (GTR) Model

A large number of Markov models have been proposed to model DNA evolution, beginning with the Jukes-Cantor model [Jukes and Cantor, 1969], which used equal mutation rates. To extend this model to allow for in-homogeneous rates, the GTR model [Tavaré, 1986] was developed. Under this model, rates of transition between states can be variable, making it more suitable for language evolution. Under language evolution, the number of states is lower than biological evolution. Traits are represented by presence or absence, whereas in biological evolution there are 4 bases when modeling DNA or RNA, and 20 bases when modeling protein sequences. The GTR model is a Markov process, consisting of a transition matrix, which is parameterised by time: t.

Let

$$P(t) = \begin{bmatrix} p_{00}(t) & p_{01}(t) \\ p_{10}(t) & p_{11}(t) \end{bmatrix} \tag{3}$$

describe the probabilities of going from one state to another within t. In the simplest case: $p_{01} = p_{10} = \mu$, the matrix is reduced to a single parameter. When applied to each site in l, the result is l'. To calculate $P(t)$, it is possible to use a *standard transition rate matrix* Q.

$$Q = \begin{bmatrix} -\sum_{i \neq j} q_{ij} & q_{ij} \\ q_{ji} & -\sum_{i \neq j} q_{ij} \end{bmatrix} \tag{4}$$

Matrix exponentiation can be used to derive $P(t)$ from Q and t:

$$P(t) = \exp(Qt) \tag{5}$$

Given $P(t)$, it becomes a relatively simple programming exercise to iterate over the sites in l, and at each site randomly pick a new state according to $P(t)$ as shown in Algorithm 1.

Algorithm 1 GTR Simulation Algorithm

```
function GTR(l, Q, T)                    ▷ Mutate l over time T at rates Q
    P = MatrixExp(QT)
    for site in l do
        site = Rand(P)                   ▷ Get a new site based on the transition
probabilities in Equation 3
    end for
    return l
end function
```

What Algorithm 1 fails to capture is the individual evolutionary events, that is, the birth and death times of cognates. If this information is desired then the algorithm can be modified. Instead of using *P(t)*, the Q matrix is used directly, as follows:

Algorithm 2 GTR Simulation Algorithm with explicit mutation events

function GTR(l, Q, T) ▷ Mutate l over time T at rates Q
    for *site* **in** l **do**
        $t \sim Exp(site\ rate)$
        while $t < T$ **do**
            $site = Rand(Q)$ ▷ Get a new site based on the transition rates in Equation 4
            $t = t + \sim Exp(site\ rate)$ ▷ Next mutation time, can be captured
        end while
    end for
    return l
end function

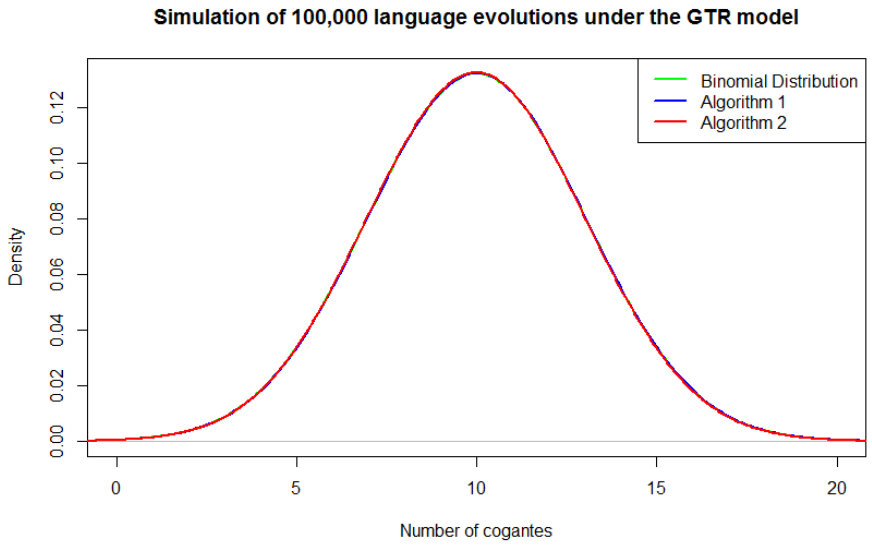

Figure 2: A comparison of the resulting languages from running Algorithms 1 and 2 100,000 times on random initial languages. The output from each algorithm was counted for the number of surviving cognates, which is plotted against the stationary binomial distribution.

Algorithm Comparison In Figure 2, a single language was independently mutated by Algorithms 1 and 2 100,000 times each. The results of the simulation show that both methods produce the same distribution of the number of cognates. Parameters used in the figure are as follows: $Q = [[-0.5, 0.5], [0.5, -0.5]]$, $T = 100$, all starting languages were 20 traits in length. The simulations are plotted with the stationary binomial distribution, where p is the probability of a 1 given a sufficiently large t – see Equation 5. A binomial distribution is used due to the binary nature of the state space, traits are either absent or present. The final parameter, n is the length of the language (20).

2.3.2 Covarion Model

The Covarion Model (where covarion is short for Concomitantly variable codons) is an extension of the standard GTR model. Instead of every site being variable, sites are able to move between variability and invariability at specific rates, δ and κ [Fitch and Markowitz, 1970]. While this was initially developed to model site-specific changes of mutation rates, it can also be used to model cognates with differing rates of evolution over time, which is more consistent with linguists' view of language evolution [Bouckaert et al., 2012].

In the Q matrix, each trait can now also move from a variant state to an invariant state, in addition to the states of trait birth and death. This means that the size of the initial matrix has to multiplied by a factor of four.

$$Q = \begin{bmatrix} -(q_{01}+\delta) & q_{01} & \delta & 0 \\ q_{10} & -(q_{10}+\delta) & 0 & \delta \\ \kappa\delta & 0 & -(\kappa\delta) & 0 \\ 0 & \kappa\delta & 0 & -(\kappa\delta) \end{bmatrix} \tag{6}$$

$$Q = \begin{bmatrix} A & B \\ C & D \end{bmatrix} \tag{7}$$

The matrix in Equation 6 can be explained as the following four quadrants in Equation 7:

1. **A** - Original Q (variant) matrix.
2. **B** - Variant to invariant rates.
3. **C** - Invariant to variant rates.
4. **D** - Mutation rate in the invariant state.

Equation 6 can then be applied to either Algorithms 1, or 2, without further modification.

2.3.3 Stochastic-Dollo Model

Under the previous models, evolutionary change across a set of languages is considered *time-reversible*. However, such an assumption is in direct contrast to the *Dollo Parsimony* assumption, under which: "An organism cannot return, even partially, to a previous state already realized in its ancestral series." [Dollo, 1893]. While many evolutionary features can be modeled using GTR it is argued that complex morphological traits are better suited to the Dollo Parsimony [Gould, 1970].

The Dollo model is useful for language evolution because the creation of new cognates is itself a complex and rare event, making it extremely unlikely the same word appears in two languages independently.

Using the simplest form of the stochastic-Dollo model (binary state), no modification is required to move from biological to lexical evolution [Nicholls and Gray, 2006]. When evolving, languages gain *unique* traits in a Poisson distributed fashion with rate λ. These traits are then lost at the *per capita* Poisson rate μ.

This results in the following stationary distribution for the number of traits in a language [Alekseyenko et al., 2008].

$$\frac{\lambda}{\mu} \tag{8}$$

Evolutionary events are Poisson distributed along a branch; with an inhomogeneous rate that combines the rates of the birth and death Poisson processes, where k is the number of traits in a language.

$$t \sim Exp(\lambda + k\mu), \tag{9}$$

Conditional on an event occurring, the probability that it is a birth event is:

$$\Pr(0 \to 1) = \frac{\lambda}{\lambda + k\mu} \tag{10}$$

And the probability that it is a death event is:

$$\Pr(1 \to 0) = \frac{k\mu}{\lambda + k\mu} \tag{11}$$

Once Equations 9, 10, and 11 have been defined for given values of λ and μ, the process can be simulated using Algorithm 3.

Algorithm 3 Stochastic-Dollo Simulation Algorithm

function SD(l, λ, μ, T) ▷ Mutate l over time T at rates λ, and μ
  $t \sim Exp(\lambda + k\mu)$
  while $t < T$ **do**
    $u \sim U(0, 1)$
    if $u \leq \frac{\lambda}{\lambda + k\mu}$ **then** ▷ Birth event
      $l += new\,Trait$
    else ▷ Death event
      $l -= Random\,Trait$
    end if
    $t += \sim exp(\lambda + k\mu)$
  end while
end function

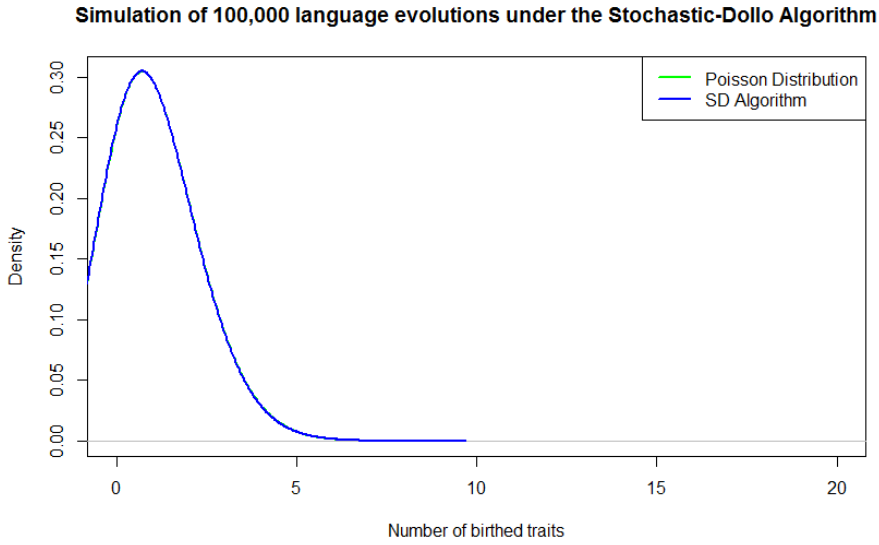

Figure 3: Result of running Algorithm 3 100,000 times, and counting the number of birthed traits. This was then plotted against the stationary Poisson distribution, where the rate is defined by Equation 8. Both λ and μ are set at 0.5, and $T = 10$.

Algorithm Validation The number of traits in the language output from Algorithm 3 are Poisson distributed, with a mean derived from Equation 8. Randomly sampling from the Poisson distribution can then be compared with the output of the algorithm to determine its correctness. Figure 3 illustrates that given sufficiently large t, any language will move towards the stationary distribution with a mean defined by Equation 8. Language l is fixed for all simulations, and is generated by giving it a number of traits drawn from the stationary distribution.

3 Whole Tree Sequence Generation

The next logical step from single lineage evolution, is to generate an entire tree from a single root language.

Once a tree structure is generated, it acts as a skeleton upon which languages are evolved. The time over which l_2 diverges from l_1 can be drawn from individual branch lengths.

Whole tree generation is an extrapolation of single lineage evolution. Each individual branch evolves according to Algorithm 1, 2, or 3.

Algorithm 4 Non-borrowing Tree Generation

```
1:  function MUTATE TREE(tree, c)        ▷ Mutate languages down tree , and build Alignment c
2:      currentParents = [tree.root]
3:      newParents = []
4:      while currentParents.length > 0 do
5:          for parent in currentParents do
6:              children = parent.getChildren
7:              for child in children do
8:                  T = parent.height − child.height        ▷ Get branch length
9:                  l = mutate(parent.lang, T)              ▷ Mutate language
10:                 child.lang = l
11:                 c.add(l)                                ▷ Add language to Alignment
12:                 newParents.add(child)
13:             end for
14:         end for
15:         currParents = newParents
16:         newParents = []
17:     end while
18:     Return tree, c
19: end function
```

Algorithm 4 iterates through the nodes of *tree*, and mutates along its branches. Line 9 is where one of the previous algorithms returns a newly evolved language for the *child* node, based on the language of its *parent*.

3.1 Algorithm Validation

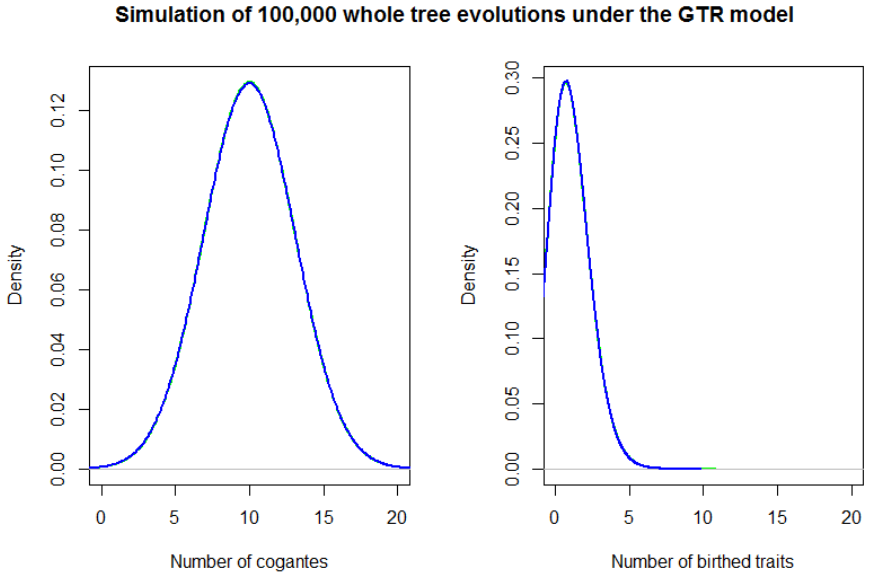

Figure 4: Comparisons of the different versions of Algorithm 4 (blue), and their underlying distributions (green) as defined in Figures 2, and 3. GTR is on the left, and SD is on the left.

Since Algorithm 4 uses previously defined algorithms, over a whole tree, the leave nodes produce the same distribution of languages as with the individual algorithms. The parameters l and T can be redefined as the root language and tree height respectively. Figure 4 demonstrates that given for random trees with a high exponential branching rate (0.005), these distributions hold true.

4 Borrowing

In language evolution, borrowing is the phenomenon wherein a cognate from one language is borrowed by another language. This is analogous to horizontal gene transfer in biological evolution. More formally, borrowing events occur from a given language randomly at a per cognate exponential rate $b\mu$. When a borrowing event occurs at time t, language l_2 receives a trait from l_1. Both the languages and traits are chosen at random, l_1 then gives a cognate to l_2, regardless of whether l_2 already has a trait or not. In the case where l_2 already has a trait, this can be considered replacement.

Borrowing runs concurrently with a cognate substitution model. In the case of GTR, borrowing is mathematically equivalent to a trait birth event, with a differing rate [Atkinson et al., 2005]. The stochastic-Dollo model complicates this slightly. Since individual traits are only born once, borrowing violates the Dollo assumption; when a trait is borrowed from l_1 to l_2, it can be considered to be born into l_2, meaning that an already dead trait can be born again.

To simulate borrowing, every evolutionary event must be captured, so as to accurately determine the number traits present in each language at any time, which determines the total rate of borrowing, and whether site i has a trait at time t, from which borrowing can occur. This means that Algorithm 1 of the GTR model cannot be used (without modification), as it doesn't explicitly capture mutation events.

4.1 Global Borrowing Model

The global borrowing model (Algorithm 5), is the simplest way of modeling borrowing. Under this model, any language can borrow from any other language at t.

Algorithm 5 Global Borrowing Simulation Algorithm

function GB(l_1, b, T) ▷ Borrow from l_1 over time T at rate b
  $t = T.height - (x \sim exp(b))$
  while $t > 0$ **do**
    $L = \{x \mid x\, is\, alive\, at\, t \wedge x \neq l_1\}$ ▷ Get all other languages L alive at t
    $l_2 = Random(x \in L)$
    $site = Random((site \in l \wedge site = 1))$
    $l_2(site) = 1$ ▷ Birth randomly selected trait in l_2
    $t + \sim exp(b)$
  end while
end function

4.2 Local Borrowing Model

Local borrowing deals with the concept that highly divergent languages are going to be less likely to borrow from one another [Nicholls and Gray, 2008]. Under local borrowing, l_2 can only borrow from l_1 if they share a common ancestor within time z. The algorithm for local borrowing is similar to Algorithm 5, and differs only in that it checks that l_1 and l_2 have a common ancestor within z. If there is not common ancestor, then the borrowing event does not occur.

Algorithm 6 Local Borrowing Simulation Algorithm

function LB(l_1, b, T, z) ▷ Borrow from l_1 over time T at rate b, with distance z
  $t = T.height - (x \sim exp(b))$
  while $t > 0$ **do**
    $L = \{x \mid x\, is\, alive\, at\, t \wedge x \neq l_1\}$ ▷ Get all other languages L alive at t
    $l_2 = Random(x \in L)$
    $site = Random((site \in l \wedge site = 1))$
    if $dist(l_1, l_2, z) == True$ **then** ▷ Check languages have common ancestor within z
      $l_2(site) = 1$ ▷ Birth randomly selected trait in l_2
    end if
    $t - \sim exp(b)$
  end while
end function

Since the rate of language borrowing changes with the number of languages alive at time t [Atkinson et al., 2005], it is not possible to use the single lineage algorithms seen previously. Instead–for both GTR and SD–new total rates of evolution have to be defined with borrowing included, and from this new algorithms can be created and validated.

4.3 Adding Borrowing to Cognate Substitution Models

4.3.1 GTR Borrowing

Under the GTR model, mutations happen at a per site rate μ, which can be represented by multiplying the mutation rate by the number of sites $|l|$, in a language l.

$$Mutation\ Rate = \mu\,|l| \tag{12}$$

Borrowing happens at a per alive (1) site rate $b\mu$, which can be represented similarly to Equation 12; where the borrowing rate is multiplied by the number of alive sites k.

$$Borrowing\ Rate = \mu b k_l \tag{13}$$

Equations 12, and 13 can be combined and extended to n languages to produce a total rate for any point on the tree.

$$Total\ Rate = \mu \left(\sum_{i=1}^{n} |l_i| \right) + b\mu \left(\sum_{i=1}^{n} k_i \right) \tag{14}$$

Once the total rate is defined, the individual probabilities can be defined in a similar vein as the original stochastic-Dollo model. It then becomes possible to construct an algorithm to traverse and populate a tree with generated languages.

Algorithm 7 GTR tree generation with borrowing

function MUTATE TREE($tree, borrow, z$) ▷ Mutate $tree$, using $borrow$ rate, and z for local borrowing distance
  $aliveNodes = \{x \mid x\ is\ alive\ at\ 0.0\}$ ▷ Get alive languages at 0.0 (root)
  $totalRate = getTotalRate(aliveNodes)$
  $t = tree.height - (x \sim exp(totalRate))$
  for $event$ in $branchTimes$ **do**
    $aliveNodes = \{x \mid x\ is\ alive\ at\ event\}$
    $totalRate = getTotalRate(aliveNodes)$
    while $t > event$ **do**
      $u \sim U(0, 1)$
      if $u < \frac{Mutation\ Rate}{Total\ Rate}$ **then** ▷ Mutation event
        $node = Random(aliveNodes)$
        $site = Random(node.language)$ ▷ Pick random site in language
        $node.language[site] = 1 - site$ ▷ $0 \rightarrow 1$ or $1 \rightarrow 0$
      else ▷ Borrowing event
        $n_1, n_2 = Random(x \in aliveNodes \wedge n_1 \neq n_2)$ ▷ n_1 and n_2 are picked randomly with probability $n_*/totalBirths(aliveNodes)$
        $site = Random((site \in n_1.language \wedge site = 1))$
        if $dist(n_1, n_2, z) == True$ **then** ▷ Check languages have common ancestor within z
          $n_2.language[site] = 1$ ▷ Birth randomly selected trait in l_2
        end if
      end if
      $totalRate = getTotalRate(aliveNodes)$
      $t - \sim exp(TotalRate)$
    end while
  end for
end function

Algorithm Validation To confirm that Algorithm 7 works, it can be compared to the multinomial distribution which underlies the interactions between multiple languages.

For two languages, each event in the tree moves the twin cognates from each language into a new state; this is defined by the following rate matrix.

$$Q = \begin{array}{c|cccc} & \mathbf{00} & \mathbf{01} & \mathbf{10} & \mathbf{11} \\ \hline \mathbf{00} & -\sum_{i \neq j} q_{ij} & \mu & \mu & 0 \\ \mathbf{01} & \mu & -\sum_{i \neq j} q_{ij} & 0 & (b+1)\mu \\ \mathbf{10} & \mu & 0 & -\sum_{i \neq j} q_{ij} & (b+1)\mu \\ \mathbf{11} & 0 & \mu & \mu & -\sum_{i \neq j} q_{ij} \end{array} \tag{15}$$

Then given predefined values for the rates ($\mu, b = 0.5$), the stationary distribution can be determined by plugging in Equation 15 into Equation 5, where t tends towards infinity.

$$\begin{array}{cccc} \mathbf{00} & \mathbf{01} & \mathbf{10} & \mathbf{11} \\ 0.22\dot{2} & 0.22\dot{2} & 0.22\dot{2} & 0.33\dot{3} \end{array} \tag{16}$$

Equation 16 then defines the probabilities for the underlying multinomial distribution. This can then be compared with the distribution of states output by a 2-leaf tree run through Algorithm 7.

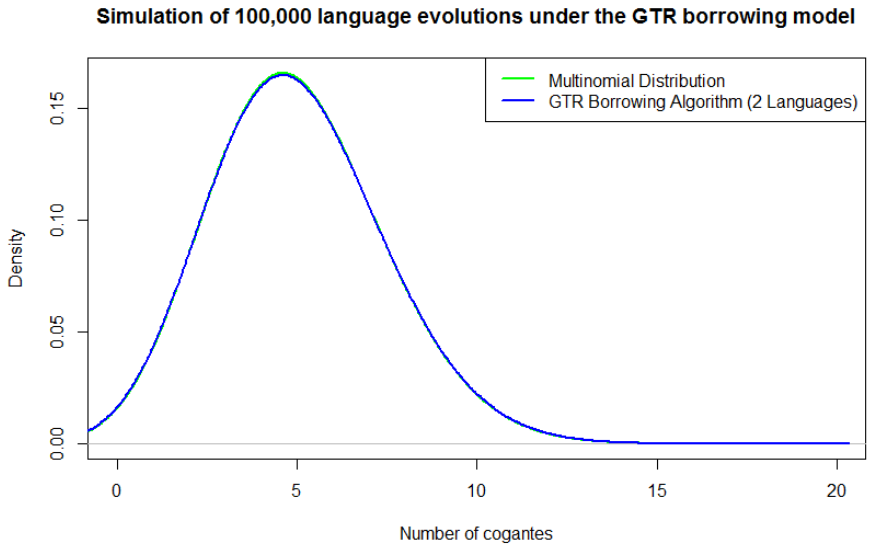

Figure 5: Comparison of theoretical multinomial distribution with an experimental multinomial distribution drawn from Algorithm 7 using two languages.

In Figure 5, there is a slight discrepancy between the two distributions. This is because the stationary distribution assumes languages of very high length;

something that is not computationally tractable for 100,000 simulations. Instead, languages were 20 cognates in length. To prove the result holds for a higher number of languages, larger matrices can be generated–from which stationary distributions can be calculated and compared.

The following calculations are for three languages.

	000	**100**	**010**	**001**	**110**	**101**	**011**	**111**
000	$-\sum_{i\neq j} q_{ij}$	μ	μ	μ	0	0	0	0
100	μ	$-\sum_{i\neq j} q_{ij}$	0	0	$(b/2+1)\mu$	$(b/2+1)\mu$	0	0
010	μ	0	$-\sum_{i\neq j} q_{ij}$	0	$(b/2+1)\mu$	0	$(b/2+1)\mu$	0
001	μ	0	0	$-\sum_{i\neq j} q_{ij}$	0	$(b/2+1)\mu$	$(b/2+1)\mu$	0
110	0	μ	μ	0	$-\sum_{i\neq j} q_{ij}$	0	0	$(b+1)\mu$
101	0	μ	0	μ	0	$-\sum_{i\neq j} q_{ij}$	0	$(b+1)\mu$
011	0	0	μ	μ	0	0	$-\sum_{i\neq j} q_{ij}$	$(b+1)\mu$
111	0	0	0	0	μ	μ	μ	$-\sum_{i\neq j} q_{ij}$

000	**100**	**010**	**001**	**110**	**101**	**011**	**111**	
0.0930	0.0930	0.0930	0.0930	0.1395	0.1395	0.1395	0.2093	(17)

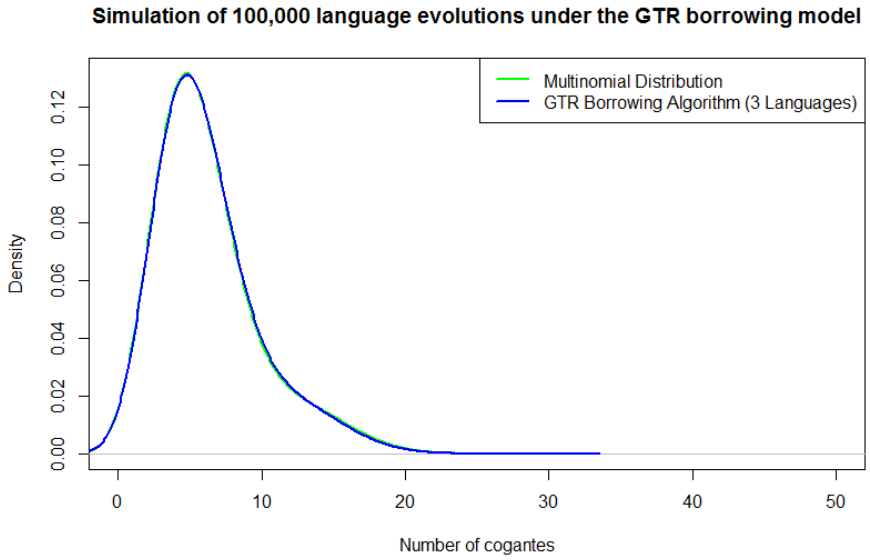

Figure 6: Comparison of theoretical multinomial distribution with an experimental multinomial distribution drawn from Algorithm 7 using three languages.

Figure 6 once again demonstrates the discrepancy between experimental data, and the theoretical distribution defined by Equation 17. However, it does demonstrate that the algorithm is effective for multiple languages.

4.3.2 Stochastic-Dollo Borrowing

Having both a birth rate — λ — and a death rate — μ – complicates the stochastic-Dollo model somewhat when borrowing is introduced. The birth rate is tied to the number of languages alive, $\left|\bar{l}\right|$, at time t.

$$Birth\ Rate = \lambda \left|\bar{l}\right| \tag{18}$$

The death rate is proportional to the number of alive cognates k in each language $i \in \left|\bar{l}\right|$.

$$Death\ Rate = \sum_{i=1}^{n} \mu k_i \tag{19}$$

Finally, the borrowing rate is related to both the death rate, and the number of alive cognates in each language.

$$Borrow\ Rate = \mu b \sum_{i=1}^{n} k_i \tag{20}$$

Combining Equations 18, 19, and 20 produces the total rate of evolution across the tree.

$$Total\ Rate = \lambda \left|\bar{l}\right| + \left(\sum_{i=1}^{n} \mu k_i\right) + \left(\mu b \sum_{i=1}^{n} k_i\right) \tag{21}$$

The algorithm for the stochastic-Dollo borrowing model is very similar to Algorithm 4, with the addition of separate birth and death events.

Algorithm 8 SD tree generation with borrowing

function MUTATE TREE($tree, c, borrow, z$) ▷ Mutate $tree$, and build Cognate Set c, using $borrow$ rate, and z for local borrowing
  $aliveNodes = \{x \mid x\ is\ alive\ at\ 0.0\}$ ▷ Get alive languages at 0.0 (root)
  $totalRate = getTotalRate(aliveNodes)$
  $t = tree.height - (x \sim exp(totalRate))$
  for $event$ in $branchTimes$ **do**
    $aliveNodes = \{x \mid x\ is\ alive\ at\ event\}$
    $totalRate = getTotalRate(aliveNodes)$
    while $t > event$ **do**
      $u \sim U(0, 1)$
      if $u < \frac{Birth\ Rate}{Total\ Rate}$ **then** ▷ Birth event
        $node = Random(aliveNodes)$
        $node.language+ = 1$ ▷ Add a cognate
      else if $u > \frac{Birth\ Rate}{Total\ Rate}$ and $u < \frac{Birth\ Rate}{Total\ Rate} + \frac{Death\ Rate}{Total\ Rate}$ **then** ▷ Death event
        $node = Random(aliveNodes)$
        $site = Random((site \in n_1.language \wedge site = 1))$
        $n_2.language[site] = 0$ ▷ Kill randomly selected trait in $node$
      else ▷ Borrowing event
        $n_1, n_2 = Random(x \in aliveNodes \wedge n_1 \neq n_2)$ ▷ n_1 and n_2 are picked randomly with probability $n_*/totalBirths(aliveNodes)$
        $site = Random((site \in n_1.language \wedge site = 1))$
        if $dist(n_1, n_2, z) == True$ **then** ▷ Check languages have common ancestor within z
          $n_2.language[site] = 1$ ▷ Birth randomly selected trait in l_2
        end if
      end if
      $totalRate = getTotalRate(aliveNodes)$
      $t + \sim exp(TotalRate)$
    end while
  end for
end function

Algorithm Validation Instead of determining the stationary distribution of three interacting Poisson processes, the algorithm can be validated against previous work. TraitLab [Nicholls et al., 2011] was created to perform some of the operations described in this thesis. Data synthesis–under the stochastic Dollo model–inside TraitLab produces sequences which can then be compared to similar output from Algorithm 8.

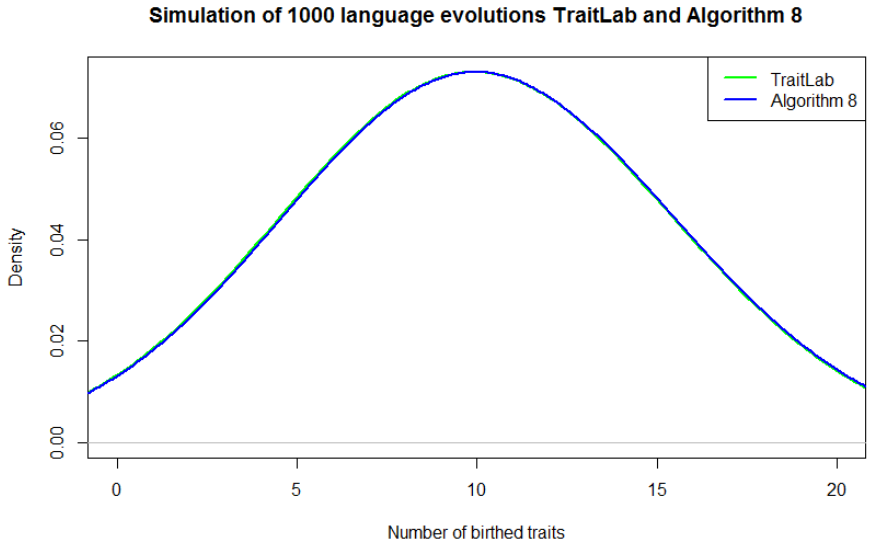

Figure 7: Comparison of TraitLab [Nicholls et al., 2011], and Algorithm 8.

Using the default trait loss rate (0.2) and borrowing rate (0.1) in TraitLab, reducing taxa to 8, and mean traits to 10, it becomes tractable to compare the two programs. To run TraitLab in batch mode, small modifications were made to the public source code by one of the authors (David Welch). Figure 7 demonstrates that both simulations produce the same distribution of languages. The number of runs is limited– comparative to previous validations–due to a limit on the number of runs TraitLab is able to do (1000).

5 No Empty Meaning Class

To keep synthesized data consistent with its real world counterpart, it is important to recognize that each language is unlikely to have no cognate for a specific meaning. This is consistent with well known word lists like the Swadesh word list, where every language has a cognate for each meaning in the list; due to the universality of the chosen words.

For this to take place, the cognate substitution models need to be modified to handle the death of a cognate in a more intelligent manner.

Algorithm 9 No Empty Meaning Class Algorithm

```
1: function DEATHCHECK(l)                ▷ Check whether a death event will reduce traits in l to 0
2:     if l.count(1) > 1 then            ▷ Check number of alive traits
3:         return TRUE                   ▷ No Empty Meaning Class isn't violated
4:     else
5:         return FALSE                  ▷ No Empty Meaning Class is violated
6:     end if
7: end function
```

Algorithm 9 is then run whenever a cognate death event occurs in prior algorithms, determining whether or not the death event is legal.

There may be instances where this behavior is not required, therefore it is coded as an optional flag.

6 Missing Data

Realistically, linguistic data is often incomplete. In the case of ancient languages the passage of history often obscures or destroys records, and through this marginally adequate data, gaps in the lexicon form. The *Comparative Indo-European Database* is one such example of a database with considerable missing data Isidore Dyen [1992].

When producing linguistic models, it is important to determine whether or not missing data is particularly impactful; in some cases encoding missing data as absent cognates produces only negligible effects on the outcome of a model Atkinson et al. [2005]. However, where missing data is more prevalent the result of such data manipulation could well be significant. This suggests that models should be validated against various types of missing data, as not to introduce biases.

Where missing data is truly random, it is still possible for that data to be representative of its population. However, where data is systematically missing in a non-random way, it can skew analysis of the data in a particular direction. For example, if half of the languages in a dataset are missing a particular cognate due to researcher error, analysis will reveal a cognate loss event somewhere in the ancestry, where no actual event exists.

6.1 Missing Languages

When a language loses its last native speaker, it becomes a *dead language*, from this point forward any words that have not been recorded are lost. Modern languages are much less likely to dodge at least partial recording due to large efforts in the linguistic community Crystal [2000]. Ancient languages are often not as safe. Hunnic is one such language where the sole evidence is a small handful of words and names Pritsak [1982]. Languages like Hunnic can still be used in a phylogenetic approach, but only if models are built to handle vast swathes of missing information.

$$C = \begin{bmatrix} 0 & 1 & 1 & 0 & 1 & 1 \\ ? & ? & 1 & ? & ? & ? \\ 1 & 1 & 1 & 0 & 0 & 1 \end{bmatrix} \tag{22}$$

Equation 22 demonstrates how a language might be included in a dataset C, absent of significant cognate information. It still may have useful metadata attached that aids in future analysis.

Algorithm 10 Missing Languages Algorithm

```
1: function MISSINGLANGUAGES(A, p)        ▷ Make languages unknown in alignment a unknown with probability p
2:     for l in A do                       ▷ For each language in a
3:         u ~ B(|l|, p)                   ▷ Get binomially distributed number of events
4:         for 0..1..u do
5:             l[r ~ U(0, |l|)] = ?        ▷ Set random cognate to missing
6:         end for
7:     end for
8: end function
```

From a technical perspective, the simplest way to randomly simulate missing languages is to do so in a binomial manner given some probability p, as demonstrated by Algorithm 10.

6.2 Missing Meaning Classes

A more common type of missing data is a missing meaning class. While word lists endeavor to be constructed from universal cognates, many cognates in these lists are not completely universal Holman et al. [2008]. Furthermore, for older languages, cognate information may be incomplete, leading to situations where models must remain accurate with fewer cognate classes that initially believed.

$$C = \left[\begin{array}{cc|cc|cc} 0 & 1 & ? & 0 & 1 & 1 \\ 0 & 1 & 1 & ? & 1 & 1 \\ 1 & 1 & 1 & ? & 0 & 1 \end{array}\right] \tag{23}$$

If Equation 23 is composed of three meaning classes of two cognates a piece, then its representative of data that is missing in the second meaning class. Algorithm 11 defines a similar binomial method, as seen in Algorithm 10.

Algorithm 11 Missing Meaning Classes Algorithm

1: **function** MISSINGMEANINGCLASSES(MC, A, p) ▷ Make meaning classes MC in alignment A unknown with probability p
2: **for** mc in MC **do** ▷ For each meaning class in MC
3: $u \sim B(|l|, p)$ ▷ Get binomially distributed number of events
4: **for** 0..1..u **do**
5: $l = rand(A)$ ▷ Get random language from A
6: $l[r \sim U(MC)] = ?$ ▷ Set random cognate in meaning class to missing
7: **end for**
8: **end for**
9: **end function**

6.3 Algorithm Validation

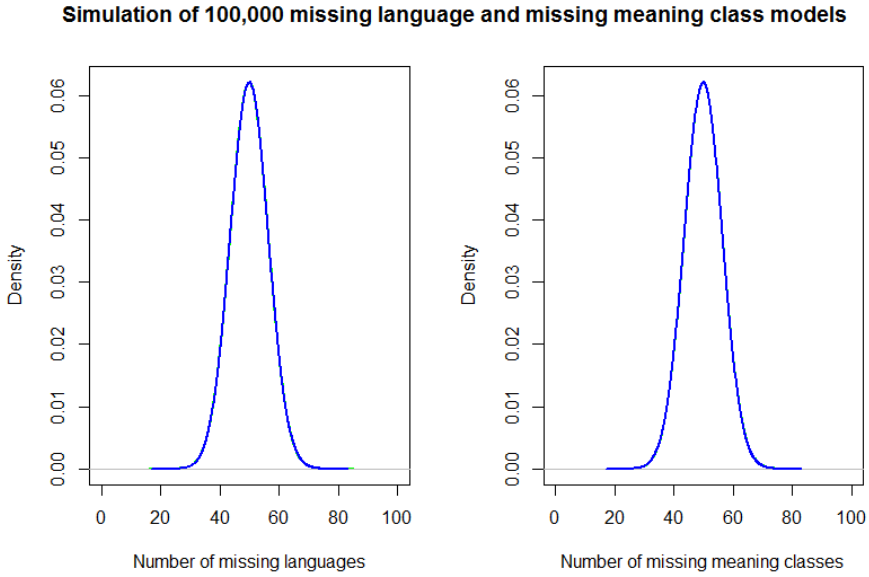

Figure 8: Comparisons of Algorithms 10 (left) and 11 (right) in blue, with their underlying binomial distributions in green.

To validate Algorithms 10 and 11, a 10 language alignment was generated, with 10 meaning classes. The algorithms were then run 100,000 with $p = 0.5$, and the resulting number of missing languages or meaning classes was plotted against the binomial distribution where $n = 10$ and $p = 0.5$.

Figure 8 demonstrates that both algorithms perform as expected, by producing the correct underlying distributions.

7 Computational Optimizations

While the algorithms described in previous sections are functionally correct, the borrowing algorithms (Algorithms 7 and 8) had to undergo additional optimizations to run in satisfactory time.

Sequence Representation Initially, sequences were represented as BEAST2 `sequence` objects, this allowed for easier interfacing with a number of BEAST2 methods that were used throughout the code. The problem with this is the algorithms perform a large amount of string manipulation, and each time a language was changed, the `sequence` object was unwrapped, modified, and wrapped up again. Obviously, this produced significant performance hits in the face of large

amounts of manipulation. To remedy this, the sequence objects were only updated at branching events, and internal primitive `String` representations were kept in the interim.

It is possible to increase the speedup further here, Java `String` objects are not ideal for manipulation. The best solution if a language is of constant length is to use a byte array. However, in the case of stochastic-Dollo model, the increasing length of the language results in a large amount of time spent copying n size arrays to $n + 1$ arrays. While this would have worked for the GTR model, code consistency was deemed more important, and speedups were found elsewhere.

Java's `StringBuilder` class was also considered, but ultimately was not implemented because it could not be determined if moving from `String` to `StringBuilder` would produce a significant gain.

Cognate Memorization After moving to a faster internal representation of languages, the next bottleneck was the calculation of cognate locations. To do most major operations (such as determining the rate of evolution), the number of alive cognates is required. This process happens in $O(n)$ time, and usually has to happen for every alive language at any singular point in time. These operations quickly add up to substantial computing time, and therefore pave the way for much needed memorization.

Instead of calculating this number on the fly, a running counter can be kept for every language, that is incrementally updated as operations on a language occur. While this does increase the memory footprint of the program, it does so with a massive increase in speed, and is certainly worth the trade-off.

Finding an Alive Cognate After optimizing the representations of the data, the next step is to find functions that are particularly time consuming. The most time consuming method–outside of the main loop–is the one to pick a random alive cognate. While the number of births have been memorized, their positions have not, leading to another computationally intensive function.

There are two ways to find a random index: deterministically, and non-deterministically. Under the deterministic method, the language indices are shuffled randomly ($O(n)$), and this shuffled list of indices is processed until an alive cognate is found ($O(n)$). The resulting method has a maximum running time of $O(2n)$, and can be beaten by the non-deterministic method under most cases.

The non-deterministic method simply randomly draws an indices from the language until it finds an alive cognate. The issue here is that the method is non-bounded and not guaranteed to complete. The other issue with this method is that under sparse structures (such as those languages created using the scholastic-Dollo model), it can perform much worse than the deterministic method. However, the non- deterministic does perform slightly better for the GTR algorithms.

To deal with the stochastic-Dollo side of the problem, memorization can once again be employed. When the number of births are captured, so can the positions, which can be incrementally updated as operations occur. However, the problem with this is that if Java `List` structures are employed, removal of indices upon death is $O(n)$, thus moving the intensive operation. The `Set` object poses a similar problem, in that randomly finding an element in a set takes $O(n)$ time; the problem here is that Java does not allow direct access to buckets, unlike C++.

To gain the speedup required, a custom set class can be defined which avoids linear time operations, but this has not currently been implemented.

8 Model Misspecification

One use of simulated data is to determine the accuracy of a models estimates under different scenarios. In the case of linguistics, some critics suggest the application of biological phylogenetics to languages and cultures isn't completely apt. It is suggested that certain language phenomena (such as word borrowing), invalidate phylogenetic assumptions; in turn, producing spurious results [Moore, 1994], and [Terrell, 1988].

8.1 Testing for Misspecification

When simulating languages, it is done down an evolutionary tree. The leaf languages of this tree can then be fed into phylogenetic inference packages based on *Markov chain Monte Carlo (MCMC)* methods to recover the original tree. The difference between the original and the tree — estimated via MCMC — is then indicative of how sensitive the inference model is to misspecification.

Given a the tree-space of all possible trees, MCMC algorithms in general randomly move through this space, proposing new trees based on the current tree then accepting and rejecting moves from one tree to another; based on the probability of the tree under the model. This process is repeated millions of times — with trees in early generations being discarded (*burn-in*) — until the MCMC algorithm reaches a stationary distribution.

Once a sufficient number of trees have been sampled by an MCMC algorithm, they must be compared to the original tree to determine the amount of misspecification. There are a number of distance metrics for such a comparison, that are discussed below.

8.2 Misspecification Due To Borrowing

The effect of word borrowing on model misspecification has been studied previously; Greenhill et al. [2010] looks specifically at the stochastic-Dollo model of language evolution and concludes that under realistic levels of borrowing, phylogenetic inference remains robust. This conclusion is consistent for both global and local borrowing. The experiment itself was performed using TraitLab, and

simulating data via BEAST2 allows for both replication of the stochastic-Dollo results, and further testing under the GTR model.

[Currie et al., 2010] modeled borrowing of cultural traits as opposed to individual words. This was done using a Brownian motion model, with the additional complication of the correlation of traits. While not an entirely similar experiment, results showed that phylogenetic inference was once again robust to realistic levels of borrowing.

One study that suggests that borrowing may cause problems is [McMahon and McMahon, 2005, pp. 111-118]. In this study, borrowing levels about 10% produced trees very different from the original. There are two possible reasons for this: firstly, borrowing only occurred on leaf nodes in the final generations, and secondly, they use less powerful distance measures to reconstruct trees. What this study does highlight however, is the necessity for studying borrowing under different experimental conditions. As well as the need for further replication where possible.

8.2.1 Material and Methods

To determine the effect of borrowing on the effect of phylogenetic estimates, languages are simulated under both the GTR, and stochastic-Dollo models of language evolution. Once a set of languages is produced (from the leaves of an input tree), BEAST2's MCMC algorithms are used to rebuild the most probable trees, which are then compared to the original input tree.

Root Languages for Sequence Generation At the root of the tree, the input language needs to represent realistic data. In the format defined in Equation 1, a root language would look like a list of alive cognates. In the specific case of the Indo- European dataset, this list is 2449 cognates in length. This root language is then subdivided into meaning classes equal to its length, representing a single trait per meaning class.

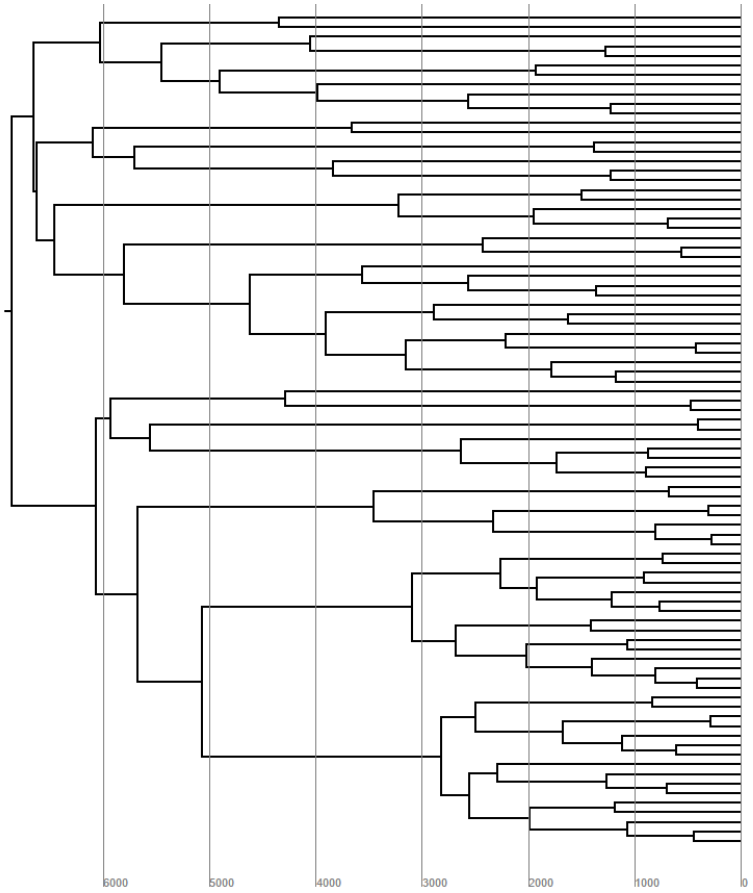

Figure 9: Tree reconstruction of Indo-European dataset for use in creating synthetic data. The tree has a root age of 6858 years.

Trees for Sequence Generation The tree itself needs to be representative of realistic data. Using the Indo-European dataset [Isidore Dyen, 1992], BEAST2 MCMC analysis was run. The result of this analysis was a sampling of probable trees, one of which was randomly chosen as the basis for the average height of generated trees. The trees themselves are randomly generated on each run, with 80 leaves, and $\lambda = 0.00055$. The tree used for the height basis can be seen in Figure 9.

Borrowing Rates To determine what rates of borrowing should be used, a relationship can be defined between the percentage of traits borrowed over a 1000 years, and the rate of borrowing b.

$$Borrowing\ Percentage = 1 - \exp(-b1000) \tag{24}$$

Table 1: Borrowing rates and associated percentages defined by Equation 24 for the GTR model.

Rate	0	0.045	0.224	0.448	0.672	0.896	1.344	1.793	2.241
Percentage	0	1	5	10	15	20	30	40	50

MCMC Run Parameters For each model, 100 MCMC runs were performed using a chain length of 22,000,000 of which 10% was discarded as burnin. However, the covarion model was run for double this length (44,000,000), as certain parameters did not converge satisfactorily in 22,000,000 steps. This resulted in 2,000 trees per run.

For all models a strict clock and yule tree priors were used. For the substitution models, GTR inference runs requires the morph-model package [Bouckaert et al., 2014a]. For Covarion inference, the Babel package was used [Bouckaert et al., 2015]. Both these packages supplied additional models not available in native BEAST2.

Tree Comparison In [Greenhill et al., 2010], the quartet distance measure is used [Day, 1986]. This measure defines the difference between two trees as the number of differing subsets of four leaves that do not have related topology. While this is a good measure, it does not take into account other tree parameters: such as tree height.

The implementation of the quartet distance measurement used is tqDist [Sand et al., 2014]. While this can't handle the output of BEAST2 directly, only small modifications are required. This distance is measured between the input tree and each inferred tree, and multiple samples per run are captured to deal with perturbations.

Height is measured as the difference between the randomly generated input tree. To account for variable starting tree heights, differences were normalised to a value between 0 and 1.

Both measures are done for a random sample of 300 trees in each MCMC run, this sample is then averaged to produce a single value per run, resulting in 100 data points per borrowing rate. In the figures below, these are plotted along with a mean for the entire rate, and the 94% highest posterior density (HPD) interval.

8.2.2 GTR Model

For the GTR model, the mutation rate μ is set such that the rate at which the probability of moving between states 1 to state 0 (or vice-versa) over a 1000 year period is 0.1, which is representative of the rate of change in the Indo-European dataset.

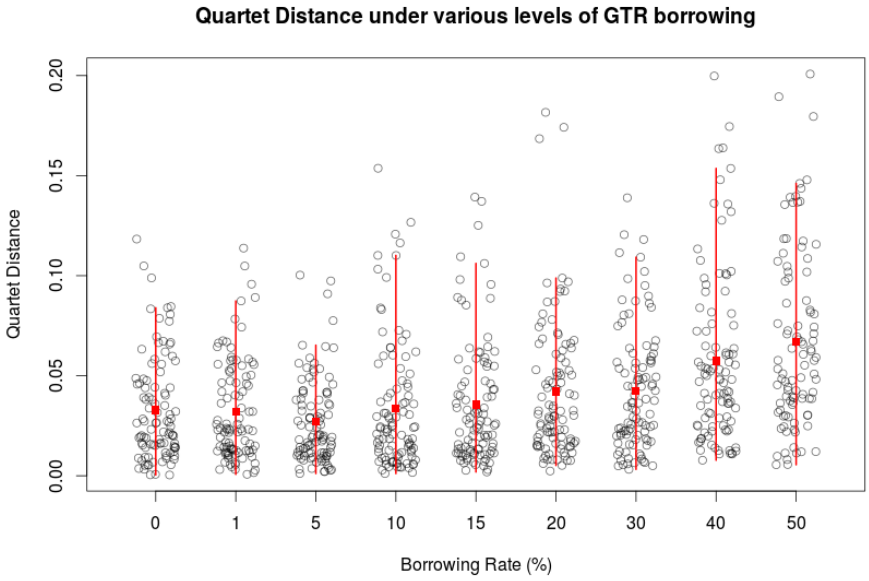

Figure 10: Quartet distances of inferred trees using synthetic languages created under the GTR model. The red square indicates the mean, and the red line indicates the 95% HPD interval.

When GTR languages are inferred under the GTR model, the mean quartet distance (Figure 10) does not deviate from the baseline (0.033), until reasonably high levels of borrowing (40% >). At 50% borrowing, the mean distance increases to 0.067.

The HPD interval begins to widen at lower levels of borrowing. Beginning at $[0.0, 0.102]$, the interval begins to increase for rates of 10% and above. At 10%, the interval increases to $[0.0, 0.115]$, and continues to increase; at 50% the interval is: $[0.0, 0.165]$.

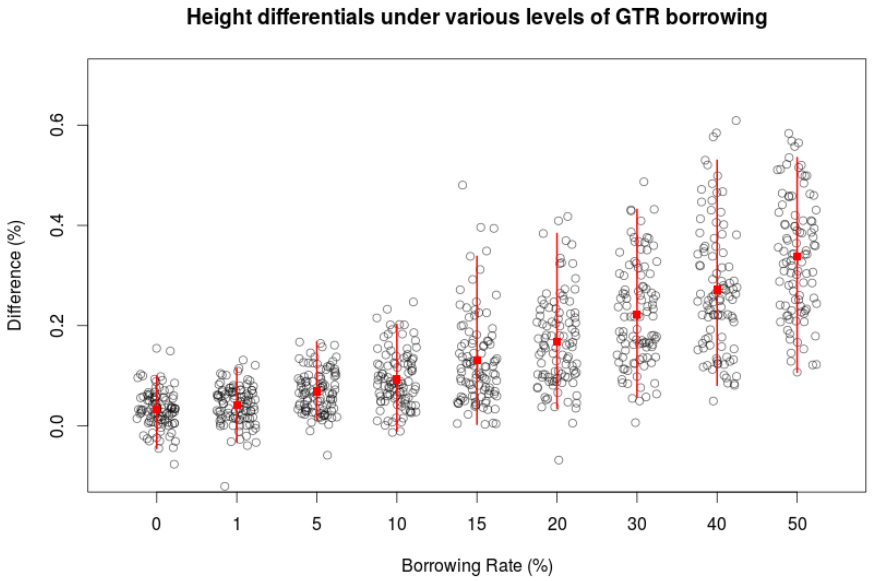

Figure 11: Percentage differences in heights between the original tree used for language generation, and the inferred trees under the GTR model. The red square indicates the mean, and the red line indicates the 95% HPD interval.

Differences between true tree height, and inferred tree height trend upward as the borrowing rate increases. Figure 11 shows that the mean height moves from 0.033 at 0%, to 0.338 at 50%. This trend begins earlier when compared to the GTR quartet distances, with the mean perceptibly increasing at the 5% level (0.068).

The HPD interval also increases with borrowing rate, starting at $[-0.106, 0.162]$ at 0% borrowing, and trending upwards to $[-0.074, 0.206]$ at 50%.

Figure 11 also highlights that a number of runs produced what could be considered outlying values.

8.3 Stochastic-Dollo Model

For the simulations, the death rate μ is defined by the mean proportion of traits lost per 1000 years. This was kept the same as the GTR model, at $l = 0.1$.

$$\mu = \frac{-\log(1-l)}{1000} \tag{25}$$

To determine the birth rate λ, the death rate is multiplied by the length–L–of the root language.

$$\lambda = L\mu \tag{26}$$

To model the stochastic-Dollo synthetic languages, two models were picked due to their common use in inference processes: GTR and Covarion.

8.3.1 GTR Inference

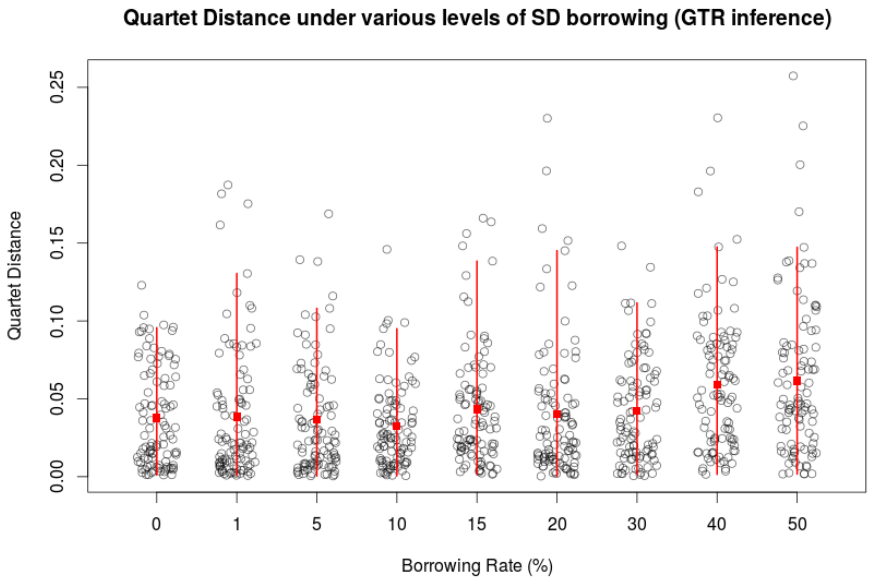

Figure 12: Quartet distances of inferred trees using synthetic languages created under the SD model, and inferred using the GTR model. The red square indicates the mean, and the red line indicates the 95% HPD interval.

When synthesizing languages under the SD model, and inferring topology under the GTR model, quartet distance increases in a similar pattern to GTR synthesis: a perceptible effect is only noticed at the highest levels of borrowing.

Deviation from the baseline (0.038) does not occur until 40% borrowing, at which point the mean increases to 0.059.

The change in HPD intervals is somewhat less clear, as the 20% interval ($[0.0, 0.149]$) is similar to the 40% interval ($[0.0, 0.147]$).

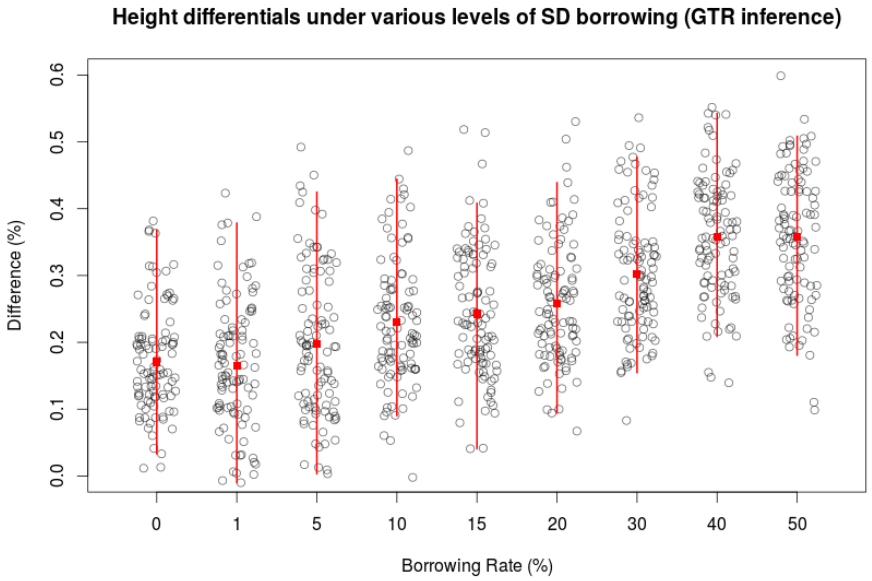

Figure 13: Percentage differences in heights between the original tree used for language generation, and the inferred trees under the SD model, and inferred under the GTR model. The red square indicates the mean, and the red line indicates the 95% HPD interval.

Figure 13 shows a strong relationship between the mean height difference and the borrowing rate. As the rate increases, the mean difference moves upwards away from the baseline (0.170) to 0.358 at 50%.

The HPD interval however, does not increase significantly. At 0% borrowing, the interval is $[-0.010, 0.372]$; at 50%, the interval is [0.160,0.541]. The difference between these two intervals is 0.1. Although, it is worth noting the upward movement of the 50% interval.

8.3.2 Covarion Inference

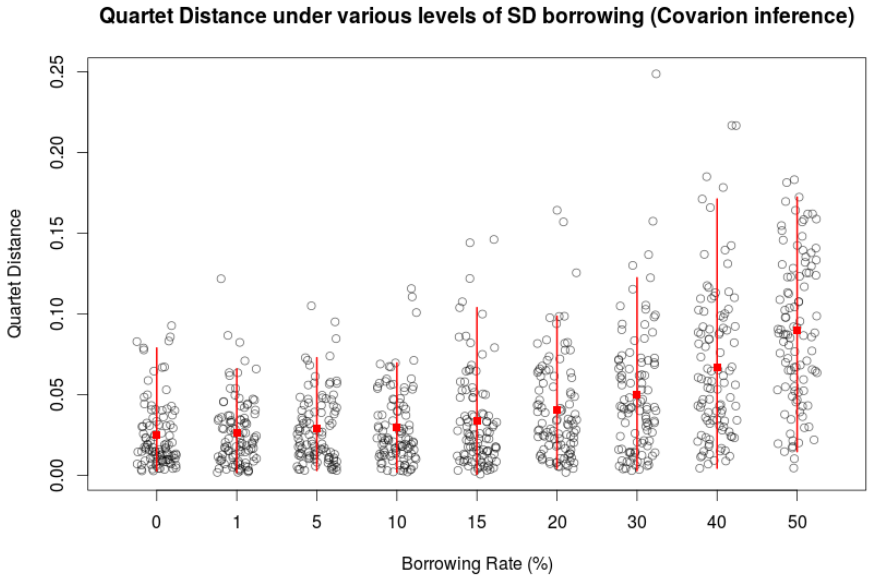

Figure 14: Quartet distances of inferred trees using synthetic languages created under the SD model, and inferred using the Covarion model. The red square indicates the mean, and the red line indicates the 95% HPD interval.

When data is inferred under the Covarion model, there is a stronger relationship between quartet distance and borrowing rate when compared to previous experiments.

In Figure 14, movement away from the mean baseline (0.025) begins at 15% borrowing (0.034), and continues to increase with the borrowing rate. At 50% borrowing, the mean quartet distance is 0.089.

This same pattern is reflected in the HPD intervals, with the baseline starting at $[0.000, 0.091]$, and 50% having an interval of $[0.000, 0.178]$.

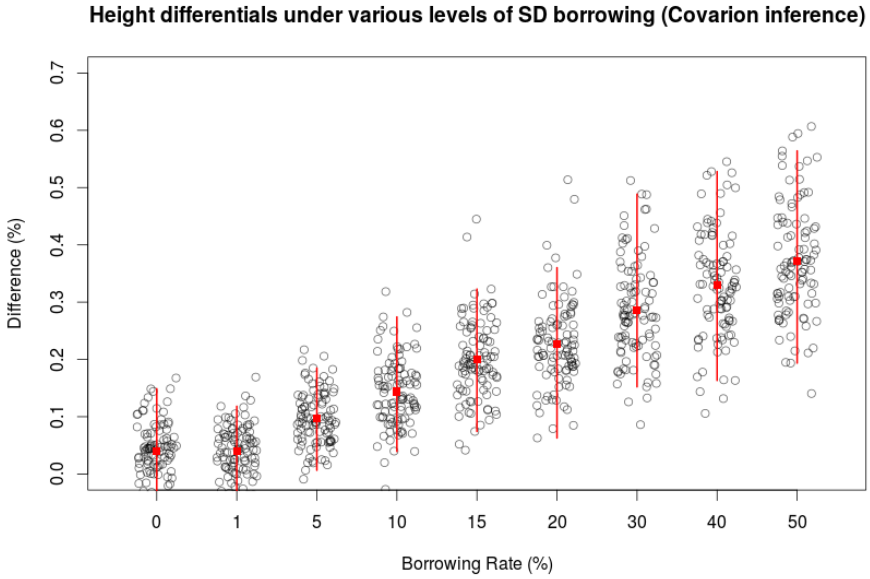

Figure 15: Percentage differences in heights between the original tree used for language generation, and the inferred trees under the SD model, and inferred under the Covarion model. The red square indicates the mean, and the red line indicates the 95% HPD interval.

Once again, the correlation between height differential and borrowing rate appears to be quite strong in Figure 15. From the mean baseline of 0.040, increases begin at 5% borrowing (0.097), and continue to 50% (0.372).

Again, this pattern is reflected in the HPD intervals. The baseline interval begins at $[-0.122, 0.197]$, and by 50% borrowing, has increased to $[0.187, 0.593]$.

8.4 Discussion

In all experiments, there is a clear correlation between an increasing borrowing rate, and further deviation from the true tree.

Before discussing the magnitude of these effects, it is important to provide context for both measurements. For quartet distance, any two random yule trees (with 80 leaves) have a mean quartet distance of 0.667, so as inferred trees approach this value, they essentially become no better than randomly assigning leaf topologies.

For the height measurement, positive differences between the true tree and inferred tree heights is symptomatic of underestimation. For example, if the true tree is 7000 years old, then a difference of 0.2, means that the inferred tree has an age of 5600 years. What this shows, is that even small height differences can actually translate into huge miscalculations.

8.4.1 Topology

For quartet distance, the relationship under both the GTR and the stochastic-Dollo model is one that doesn't take serious effect until very high levels of borrowing. All three simulations show a similar range of quartet distances ($[0, 0.25]$), and both the GTR and SD (Covarion inferred) simulations show an increasing mean quartet distance from 15% borrowing. However, increased variability in the SD (GTR inferred) data means that this relationship is less clear until around 40%.

The SD (GTR inferred) distances have a different relationship when compared to the SD (Covarion inferred) distances. In Figure 12, the relationship is relatively unclear until about 40% borrowing. However, Figure 14 shows a much clearer relationship beginning as early as 15%. This is because while the Covarion model begins with a lower mean (0.025 versus 0.038), by the end of the graph, its mean is in fact higher (0.089 versus 0.067). What this means is that SD data is better inferred under the Covarion model, as there are fewer outliers in the Covarion case.

What is important to remember that while there are obvious relationships, the mean distances are very small for all levels of borrowing. Where inference is only concerned with tree topology, reasonable borrowing appears to have a minimal effect.

This isn't to say that there isn't cause for concern at high levels of borrowing. Both the GTR and the SD have multiple outliers significantly higher than the baseline of 0% borrowing. While GTR only reaches 0.311, the SD model (inferred under both the GTR and Covarion models) has outliers as high as 0.686, which is above what is considered a random quartet distance.

8.4.2 Tree Height

When attempting to infer heights under levels of borrowing, the effects are much more pronounced. Under the GTR simulation, a steady increase to a mean difference of 0.338 (50%) results in the average age of a 7000 year old tree being underestimated by 2336 years. This effect is intensified by a steadily increasing HPD interval, which highlights the uncertainty of measurements.

Even at smaller (15%) borrowing levels, a 0.131 mean difference results in underestimation by 917 years. What is also noticeable, is that even at this level the range is very large, with outliers as far out as 0.762 (underestimation of 5334 years). Both the GTR and SD (Covarion inferred) simulations show clear trends of increasing underestimation as the borrowing rate increases.

The relationship is less pronounced in the SD model (GTR inferred), but only because the baseline mean is already significantly underestimating true tree height by 0.167. At the highest levels of borrowing, the mean and outlying estimations are similar to the GTR model, but the the 15% level has a markedly higher mean of 0.236. Increased scatter in the SD (GTR inferred) simulation also highlights that Covarion is a better inference method, as stated in the previous section.

When performing inference under the Covarion model we see a strong relationship between borrowing level and inferred tree height. Once again, this due to the fact that the Covarion baseline mean begins lower than the SD (GTR inferred) mean (0.04 versus 0.170). However, as borrowing increases the difference between the two inference methods decreases, and both end up with a similar mean at the 50% level (0.372 versus 0.358).

Ultimately, trying to infer tree height under some level of borrowing is going to result in the underestimation of the date, increasing in severity as the borrowing rate increases.

9 Conclusion

Due to extensive work already completed by the group behind BEAST2, extending it to synthesize languages is a more surmountable–but still difficult–challenge. Since one of the primary drivers of generating synthetic data is model validation, being able to generate and then experiment on models in a single system is a real advantage over previous methods.

While the basis for language generation models has been created in this thesis, it has been done in such a way that new models of generation can be implemented as needed by researchers, without having to write entire new plugins for BEAST2, or new external systems.

BEAST2 did have sequence generation already built in. However, this generation was aimed at biological data, and did not have features specific to linguistic models, such as borrowing and missing cognate models. These models are once again built in such a way that they can be extended and altered in ways that any future analyses may demand.

Finally, as a proof of work, the effects of linguistic borrowing on inference were studied. The results of this analysis not only demonstrate the effectiveness of the tools described previously, but highlight the borrowing effect on both the topology and heights of inferred trees. As the borrowing rate increases, both quartet distance and height difference increase, resulting in lower accuracy of inferred trees. Understanding and accounting for these borrowing effects will allow for future work to produce more accurate descriptions of language relationships.

References

A. V. Alekseyenko, C. J. Lee, and M. A. Suchard. Wagner and dollo: A stochastic duet by composing two parsimonious solos. Systematic Biology, 57(5): 772–784, 2008. ISSN 10635157. doi: 10.1080/10635150802434394.

Q. Atkinson, G. Nicholls, D. Welch, and R. Gray. From words to dates: water into wine, mathemagic or phylogenetic inference? Transactions of the Philological Society, 103(2):193–219, 2005. ISSN 00791636. doi: 10.1111/j.1467-968X.2005.00151.x.

K. Bergsland and H. Vogt. On the validity of glottochronology. Current Anthropology, 3(2):pp. 115–153, 1962. ISSN 00113204. URL http://www.jstor.org/stable/2739527.

R. Bouckaert, P. Lemey, M. Dunn, S. J. Greenhill, A. V. Alekseyenko, A. J. Drummond, R. D. Gray, M. A. Suchard, and Q. D. Atkinson. Mapping the origins and expansion of the indo-european language family. Science, 337 (6097):957–960, 2012.

R. Bouckaert, A. Gavryushkina, A. Drummond, and J. Heled. morph-models. https://github.com/CompEvol/morph-models, 2014a.

R. Bouckaert, J. Heled, D. Khnert, T. Vaughan, C.-H. Wu, D. Xie, M. A. Suchard, A. Rambaut, and A. J. Drummond. Beast 2: A software platform for bayesian evolutionary analysis. PLoS Comput Biol, 10(4):e1003537, 04 2014b. doi: 10.1371/journal.pcbi.1003537. URL http://dx.doi.org/10.1371%2Fjournal.pcbi.1003537.

R. Bouckaert, T. Vaughan, and S. J. Greenhill. Babel: Beast analysis backing effective linguistics. https://github.com/rbouckaert/Babel, 2015.

N. E. Collinge, E. KOERNER, and R. ASHER. The main strands of 19th century linguistics: History of comparative linguistics. Concise History of Language Sciences: From the Sumerians to the Cognitivists, 1995.

D. Crystal. Language death. Ernst Klett Sprachen, 2000.

T. E. Currie, S. J. Greenhill, and R. Mace. Is horizontal transmission really a problem for phylogenetic comparative methods? a simulation study using continuous cultural traits. Philosophical Transactions of the Royal Society of London B: Biological Sciences, 365(1559):3903–3912, 2010. ISSN 0962-8436. doi: 10.1098/rstb.2010.0014. URL http://rstb.royalsocietypublishing.org/content/365/1559/3903.

C. Darwin. The origin of species by means of natural selection. 1859.

W. H. Day. Analysis of quartet dissimilarity measures between undirected phylogenetic trees. Systematic Biology, 35(3):325–333, 1986.

L. Dollo. The laws of evolution. Société belge de géologie, 1893.

A. J. Drummond, M. A. Suchard, D. Xie, and A. Rambaut. Bayesian phylogenetics with beauti and the beast 1.7. Molecular biology and evolution, 29(8): 1969–1973, 2012.

W. Fitch and E. Markowitz. An improved method for determining codon variability in a gene and its application to the rate of fixation of mutations in evolution. Biochemical Genetics, 4(5):579–593, 1970. ISSN 0006-2928. doi: 10.1007/BF00486096. URL http://dx.doi.org/10.1007/BF00486096.

S. Gould. Dollo on dollo's law: Irreversibility and the status of evolutionary laws. Journal of the History of Biology, 3(2):189–212, 1970. ISSN 0022-5010. doi: 10.1007/BF00137351. URL http://dx.doi.org/10.1007/BF00137351.

R. D. Gray and F. M. Jordan. Language trees support the express-train sequence of austronesian expansion. Nature, 405(6790):1052, 2000. ISSN 00280836. URL http://ezproxy.auckland.ac.nz/login?url=http://search.ebscohost.com/login.aspx?direct=true&db=aph&AN=3307969&site=ehost-live&scope=site.

R. D. A. Q. D. Gray. Language-tree divergence times support the anatolian theory of indo-european origin. Nature, 426(6965):435–439, 2003. ISSN 00280836. doi: 10.1038/nature02029.

S. J. Greenhill, Q. D. Atkinson, A. Meade, and R. D. Gray. The shape and tempo of language evolution. Proceedings of the Royal Society of London B: Biological Sciences, 277(1693):2443–2450, 2010. ISSN 0962-8452. doi: 10.1098/rspb.2010.0051.

E. W. Holman, S. Wichmann, C. H. Brown, V. Velupillai, A. Müller, and D. Bakker. Explorations in automated language classification. Folia Linguistica, 42(3-4):331–354, 2008.

P. B. Isidore Dyen, Joseph B. Kruskal. An indoeuropean classification: A lexicostatistical experiment. Transactions of the American Philosophical Society, 82(5):iii–132, 1992. ISSN 00659746. URL http://www.jstor.org/stable/1006517.

T. H. Jukes and C. R. Cantor. Evolution of Protein Molecules. Academy Press, 1969.

J. B. Kruskal, I. Dyen, and P. Black. Some results from the vocabulary method of reconstructing language trees. Lexicostatistics in genetic linguistics, pages 30–55, 1973.

A. McMahon and R. McMahon. Language Classification by Numbers. Oxford University Press, UK, 2005.

J. H. Moore. Putting anthropology back together again: the ethnogenetic critique of cladistic theory. American Anthropologist, pages 925–948, 1994.

G. K. Nicholls and R. D. Gray. Quantifying uncertainty in a stochastic model of vocabulary evolution. Phylogenetic methods and the prehistory of languages Forster P., Renfrew C, pages 161–171, 2006.

G. K. Nicholls and R. D. Gray. Dated ancestral trees from binary trait data and their application to the diversification of languages. Journal of the Royal Statistical Society: Series B (Statistical Methodology), 70(3):545–566, 2008. ISSN 1467-9868. doi: 10.1111/j.1467-9868.2007.00648.x.

G. K. Nicholls, R. Ryder, and D. Welch. Traitlab: A matlab package for fitting and simulating binary tree-like data. 2011.

O. Pritsak. The hunnic language of the attila clan. Harvard Ukrainian Studies, 6(4):428–476, 1982. ISSN 03635570. URL http://www.jstor.org/stable/41036005.

A. Sand, M. K. Holt, J. Johansen, R. Fagerberg, G. S. Brodal, T. Mailund, and C. N. S. Pedersen. tqdist: A library for computing the quartet and triplet distances between binary or general trees. BMC Bioinformatics, yy(xx):ii–jj, 2014. ISSN 1471-2105.

R. D. Sleator. Phylogenetics. Archives of Microbiology, 193(4):235–9, 04 2011. URL http://ezproxy.auckland.ac.nz/login?url=http://search.proquest.com/docview/856945306?accountid=8424. Copyright - Springer-Verlag 2011; Last updated - 2014-04-02.

M. Swadesh. Towards greater accuracy in lexicostatistic dating. International Journal of American Linguistics, 21(2):pp. 121–137, 1955. ISSN 00207071. URL http://www.jstor.org/stable/1263939.

S. Tavaré. Some Probabilistic and Statistical Problems in the Analysis of DNA Sequences, volume 17, pages 57–86. Amer Mathematical Society, 1986. ISBN 0821811673. doi: citeulike-article-id:4801403.

J. Terrell. History as a family tree, history as an entangled bank: constructing images and interpretations of prehistory in the south pacific. Antiquity, 62:642–657, 12 1988. ISSN 1745-1744. doi: 10.1017/S0003598X00075049. URL http://journals.cambridge.org.ezproxy.auckland.ac.nz/article_S0003598X00075049.

A Language Sequence Generation Plugin Guide

The plugin described in this thesis can be found at: `https://github.com/lutrasdebtra/Beast-Borrowing-Plugin`.

A.1 Quick Start

The project can easily be cloned as a normal java project, however, it does require the BEAST2 [Bouckaert et al., 2014b] project in it's build path, a copy of which is included in the `lib` folder. Alternately, it can run directly as a `jar`, either through compilation, or from `Jars/LangSeqGen.jar`.

A.1.1 Command Line

Like the original BEAST2 `seqgen`, this plugin uses the same format for command line runs:

```
java LanguageSequenceGen <beast file> <nr of
    instantiations> [<output file>]
```

- The `<beast file>` is an `xml` file that specifics the initial input parameters. An example is provided below.
- To determine the number of meaning classes, `<nr of instantiations>` is provided, the position of first cognate in each meaning class is provided as an additional sequence at the end.
- If an `<output file>` is not provided, the output will be written to `std.out`.

A.1.2 BeastBorrowingPluginTest

Like most BEAST2 plugins, this plugin has its own testing suite defined in `BeastBorrowingPluginTest`. In this class, the `SeqGenTest()` runs the plugin using arguments defined within the function:

```
private static void SeqGenTest() {
        String[] args ={"examples/testSeqLangGen.xml",
        "2",
        "examples/output.xml"};
        LanguageSequenceGen.main(args);
}
```

The format of the arguments are the same as those in the *Command Line* section.

There are a number of other tests in the class that produce `csv` files, which are in turn used to validate various portions of the plugin in R.

A.2 Explanation of the Input/Output Files

A.2.1 The BEAST File

The BEAST file outlines how to produce the synthetic data. An example is provided in `/examples/testSeqLangGen.xml`; it is reproduced below:

```
<beast version='2.0'
       namespace='beast.evolution.alignment:
       beast.evolution.substitutionmodel'>

    <tree id='tree' spec='beast.util.TreeParser'
       IsLabelledNewick='true'
       newick='((((english:0.02,(german:0.01,
    french:0.01):0.4):0.01,spanish:0.3):0.2,
    italian:0.6):0.3,irish:0.6)'/>


    <run spec="beast.app.seqgen.LanguageSequenceGen"
       tree='@tree'>
        <root spec='Sequence'
           value="01010101010100100010101010000100"
           taxon="root"/>

        <subModel
           spec='ExplicitBinaryStochasticDollo'
           birth="0.5" death = "0.5" borrowrate
           ="0.0" borrowzrate="0.0"
           noEmptyTrait="false"/>
    <missingModel spec='MissingLanguageModel'
       rate="0.5"/>
    </run>
</beast>
```

- The `tree` takes a newick formatted tree with both branch distances and taxon node names.
- The `run` initiates the plugin using the `tree` defined above. It also has a number of interior parameters:
  - `root` is the sequence to be placed at the root of the tree. It should consist of present (1) or absent traits (0). The plugin does not handle missing or unknown traits. The `taxon` does not need to be *root*.

- `subModel` defines the model used to simulate evolution down the tree. All models have a `borrowrate` parameter, which defines the rate of global borrowing; `borrowzrate` defines the distance of local borrowing; note: if `borrowzrate` is set to `0.0`, the plugin assumes an infinite distance. Currently there are two models:
  - `ExplicitBinaryGTR` evolves the `root` via a Generalised Time-Reversible model. This model has a single `rate` parameter which defines the rate at which traits both can be birthed, and die.
  - `ExplicitBinaryStochasticDollo` evolves the `root` via a stochastic-Dollo model of sequence evolution, which has both a `birth` rate of traits, and a separate `death` rate.
- `missingModel` defines the model used to simulate missing data in the final alignment. Currently, this is non-optional and to not use it `rate` should be set to `0`.
  - `MissingLanguageModel` - Each language has a random binomial number of missing events, which convert random cognates in the language to `?`.
  - `MissingMeaningClassModel` - Each meaning class has a random binomial number of missing events, which convert random cognates in the meaning class to `?`.

A.2.2 The Output File

The Output file is a simple BEAST2 `alignment` piped to `xml`. An example from `/examples/output.xml` can be found below:

```
<beast version='2.0'>
<data id='SD' dataType='binary'>
    <sequence taxon='english'
        value='111111111111111111111111'/>

    <sequence taxon='german'
        value='111111111111111111111111'/>

    <sequence taxon='french'
        value='111111111111111111111111'/>

    <sequence taxon='spanish'
        value='111111111111111111111111'/>

    <sequence taxon='italian'
        value='111111111111111111011111'/>

    <sequence taxon='irish'
        value='111111111111111101101111'/>
</data>

<!-- Meaning Classes: 0 -->
<!-- Created at: 2016-04-26 15:09:16.506 -->
</beast>
```

A.3 The Thesis Analysis Package

Included in the repository is a number of classes inside `beastborrowingplugin/thesisanalysis`. These classes are used to do batch analysis of sythetic languages produced by this package under BEAST2 inference.

These classes are not required for the main running of the program, but may be useful if batch analysis is needed.